arXiv:1409.1253v1 [physics.comp-ph] 3 Sep 2014

Kinetic Activation-Relaxation Technique and Self-Evolving Atomistic Kinetic Monte Carlo: Comparison of on-the-fly kinetic Monte Carlo algorithms

Laurent Karim Béland, Yury N. Osetsky, Roger Stoller

Materials Science and Technology Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831-6138, USA

Haixuan Xu*

Department of Materials Science, University of Tennessee, Knoxville,TN,USA

Abstract

We present a comparison of the kinetic Activation-Relaxation Technique (k-ART) and the Self-Evolving Atomistic Kinetic Monte Carlo (SEAKMC), two off-lattice, on-the-fly kinetic Monte Carlo (KMC) techniques that were recently used to solve several materials science problems. We show that if the initial displacements are localized the dimer method and the Activation-Relaxation Technique *nouveau* provide similar performance. We also show that k-ART and SEAKMC, although based on different approximations, are in agreement with each other, as demonstrated by the examples of 50 vacancies in a 1950-atom Fe box and of interstitial loops in 16000-atom boxes. Generally speaking, k-ART's treatment of geometry and flickers is more flexible, e.g. it can handle amorphous systems, and rigorous than SEAKMC's, while the later's concept of active volumes permits a significant speedup of simulations for the systems under consideration and therefore allows investigations of processes requiring large systems that are not accessible if not localizing calculations.

Keywords: off-lattice kinetic Monte Carlo, Iron, saddle-search, vacancy

*Corresponding author
Email address: xhx@utk.edu (Haixuan Xu)

1. Introduction

On-the-fly kinetic Monte Carlo (KMC) methods are increasingly popular numerical tools to study diverse structures and processes in materials science. Some examples include defect diffusion in SiC [1], Fe [2, 3, 4, 5, 6, 7, 8, 9, 10], Ag [11, 12, 13], Pt [14], Al [15], Cu [16, 17], Pd on MgO [18], methanol on Cu [19], CO on ice [20], binary alloys [21], *c*-Si [22, 2, 3] and *a*-Si [2, 3, 23, 24], as well as other more complex processes such as surface erosion, thin film growth [25], sputtering of Au surfaces [26] and post-ion bombardement annealing in *c*-Si [27, 28]. This enthusiasm is justified by these methods ability to handle off-lattice positions, executing high-energy barriers with ease at any temperature (thus potentially reaching long simulation times), while requiring little *a priori* knowledge of the nature of the physical processes involved in the kinetics.

Sampling saddle points surrounding any given minimum in the potential energy landscape very efficiently is crucial to the success of such an endeavor. Indeed, since the list of processes governing the kinetics is built on-the-fly, these methods must quickly evaluate all the relevant exit states linked to a given configuration. Some methods, such as the Self-Evolving Atomistic KMC [6, 7] (SEAKMC), sample a relatively small number of such transition states each step (several tens or hundreds), while self-learning methods, such as the kinetic Activation Relaxation Technique [22, 2, 23] (k-ART), sample a larger number of states when encountering a configuration never observed in the past, while sampling a smaller number when encountering a configuration that was previously visited (adding events stored in memory to the list of processes that must be taken into account). In all cases, one can hardly understate the importance of efficiently finding saddle points.

Current methods model the kinetics using harmonic transition state theory [29, 30, 31] (hTST), choosing a constant reaction rate pre-factor. In this context, saddle searches are conducted using either minimum-mode following methods or

high temperature molecular dynamics combined with the Nudged Elastic Band [32, 9] method. The latter requires prior knowledge to be able to detect when the temperature assisted MD [33] has crossed a saddle point. The minimum-mode following methods have been more widely used in the on-the-fly KMC context. Thus, this article will focus on these.

Minimum-mode following methods include the Dimer method [34], Raleigh-Ritz minimization [35], the hybrid eigenvector following method [36] and the activation relaxation technique *nouveau* [37, 38] (ARTn). While some comparisons between these techniques have been made in the past [39, 40, 41], prescriptions concerning the optimal choice of method for studying a given system are lacking. Indeed, most problems of interest have energy landscape that can significantly differ from the idealized potential energy landscapes of the simple defects that are used when designing these numerical methods.

In addition to the choice of saddle-search method, the on-the-fly kMC techniques use different sets of approximations in order to accelerate the simulations. While molecular dynamics (MD) [5, 8], and experiments [27] indicate that these simulations captured the relevant physics, there are few direct comparison of these accelerated methods, which could assess the validity of these approximations. Such a comparison has been attempted in a model-system: the aggregation of 50 vacancies in bcc-Fe. This system was simulated using Autonomous Basin Climbing (ABC) [42], k-ART [4, 10], SEAKMC [5] and Molecular-Dynamics Saddle-Search Adaptive KMC (MDSS-aKMC) [9]. At 323K, all 50 vacancies in a 2000-atom box were found by SEAKMC to aggregate into clusters within 0.01 ms. k-ART estimated this time to be 0.1ms, while ABC estimated it at about 1 day. MDSS-aKMC simulated this process at 423K and found an aggregation time of roughly 1 ms. It was suggested that ABC's large estimate was related to the use of a single barrier, which under-samples the system's transition states[4]. While the k-ART and SEAKMC simulations where performed using very similar parameters, their respective predictions vary by one order of magnitude. To our knowledge, while there were attempts to establish the accuracy of this class of methods [43, 44, 9], no other comparison of

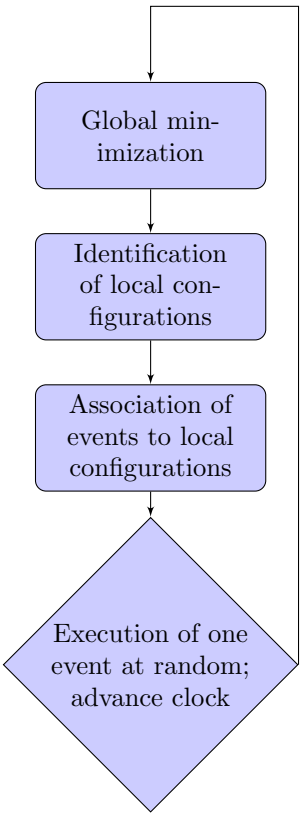

Figure 1: Simplified workflow for on-the-fly kMC algorithms, such as k-ART and SEAKMC.

such accelerated methods exists.

In this article, we first compare the two most widely methods used in modern on-the-fly kMC, the dimer method and ARTn. We concentrate on four case studies: the Fe mono-vacancy, the Fe self-intersitial atom, the Fe 50-vacancy system and FeCr dislocation loops. Secondly, we perform a systematic comparison of k-ART and SEAKMC for the 50-vacancy problem and for the transformation of two $1/2 < 111 >$ interstitial-loops in Fe into single $< 100 >$ or $1/2 < 111 >$ loops.

2. Methods

2.1. Similarities and Differences of k-ART and SEAKMC

k-ART and SEAKMC follow a similar workflow, which is illustrated in Fig. 1. The four steps described in the chart are the basis of kMC: global minimization, identification of local configurations, association of events to local configurations and the execution of one event at random (advancing the clock). For each of these steps, the two methods use a different set of approximations to accelerate simulations. These are summarized in table 1. More detailed descriptions of these algorithms can be found in numerous references [22, 2, 6, 23, 3, 7]. The global minimization is done using FIRE [45] in k-ART and conjugate gradient in SEAKMC. This step is not a bottleneck for the system sizes that we are studying (2000 to 20000 atoms).

k-ART and SEAKMC identify local configurations differently. k-ART uses a automorphism identification software package, NAUTY [46, 47], to identify connectivity graphs formed by an atom's neighboring atoms[22]. This procedure is lightweight, only requiring a small fraction of the overall computational effort. Event generation will be concentrated on atomic neighborhoods which topology does not correspond to the crystal. One should note that while a cut-off is used to determine topology, all saddle-searches are free to displace all atoms in the system. Furthermore, event search for previously visited topologies will not be attempted from scratch, but will rather recycle and refine events that were found in the past. Meanwhile, SEAKMC simply identifies defects by comparing the atomic positions of atoms to those of a perfect lattice or can be based on energy or stress criteria. The topology-based approach is more flexible and demands less *a priori* knowledge of the system. On the other hand, by precisely identifying defects, one can make saddle-searches more effective by centering them on these positions, which can increase the success rate of finding interesting events.

k-ART and SEAKMC also differ when associating events to local configurations. The former takes full advantage of the topology analysis to store and catalog events for latter use. Indeed, events are generated from scratch when

new topologies are discovered, and are recycled and refined when these topologies are revisited. Generation of completely new events for known topologies is also done periodically, going as $log_{10}(n)$, where n is the number of times a topology is encountered during a KMC run. When recycling and refining, these events associated with topologies are mapped to specific atomic configurations through geometric transformations that make use of the canonical representation of graphs generated by NAUTY. While this procedure works well for many classes of events, the canonical representation of atomic neighborhoods may leave ambiguities concerning the reconstruction of events, resulting in bad reconstruction of events. In these cases, configurations are still refined to the nearest saddle-point, which may not be related to the original event, but that still has physical meaning. The correspondance of topology to geometry is one of the main approximations of k-ART. Also, in order for k-ART to produce good kinetics, most of the important transition states must be sampled when first encountering a topology, since it will not be resampled often in the future. This is particularly important when using the basin auto-constructing Mean-Rate Method (bacMRM) [48, 2] to handle flickers, as the states in the basin will not be revisited.

In the case of SEAKMC, the main approximation is the concept of active volumes. The algorithm identifies a region surrounding each defect where event searches are constrained. This significantly reduces the computational cost of force evaluations and generally makes saddle-search and transition state sampling more efficient [6, 49, 7, 50], at the risk of neglecting some very long-range interactions, since the radius of active volumes are at least a few lattice constants. Furthermore, investigation of single vacancy diffusion in bcc Iron[5] shows that adding duplicate saddle points to the catalog does not change the defect dynamics or randomness of the diffusion. Therefore, SEAKMC sometime does not verify that duplicate events are associated with the same defect, which can lead to inaccuracies in the total rate calculation. However, while errors in multiplicity affect rates linearly, this effect is usually small compared to errors on barrier height, that effect rates exponentially. Also, SEAKMC does not garan-

tee detailed balance, i.e. it does not assure that the events necessary to reverse a trajectory are included in the catalog. While this formally excludes proper simulations at thermodynamic equilibrium, kinetics of off-equilibrium systems should be correctly captured. In principle, with more saddle point sampling, this issue could be resolved. However, for complex defects, obtaining the required sampling is challenging. Finally, SEAKMC will deliberately sample an incomplete catalog of events for each defect at each KMC step, correcting for this by flushing this catalog when an event is executed and resampling a new uncorrelated catalog.

In addition, both of these methods are based on harmonic Transition State Theory (hTST) and constant pre-factors for transition rates. However, these two approximations are not inherent to these off-lattice KMCs. In principle, one could calculate pre-factors using harmonic or quasi-harmonic theories. One could even perform free energy calculations of these barriers, albeit at a great computational cost. When executing events, both methods also treat transitions as uncorrelated Poisson processes, the basic approximation of kMC. Moreover, k-ART approximates super-basins of states connected by small barriers (flickers) by postulating that the mean-passage time adequately captures the relevant physics. This approximation is implemented through the bacMRM.

2.2. Set-up of Saddle-search Methods Comparison

In order to compare the dimer method and ARTn, we implement both in the Self-evolving atomistic kinetic Monte Carlo (SEAKMC). One of SEAKMC's main features is the use of the active volume concept, which are local regions encompassing the defects in the system of interest that limit the region where atoms will be displaced during the saddle search. This concept is very useful in mostly crystalline systems where defects can be readily identified. By limiting the region where force calculations need to be performed, SEAKMC can execute saddle searches much faster than in the case where force calculations are performed across the whole system. By ensuring that both the dimer method and ARTn perform searches constrained to the same subset of atoms and the

same implementation of the potentials, we try to make this comparison as fair as possible. Also, we adjusted the Dimer and ARTn search parameters to make them as efficient as possible for each system.

We use a recent implementation of ARTn, which uses improvements described in several studies [51, 52, 41, 23]. However, we do not use the Direct Inversion in the Iterative Subspace [53], but instead use the Newton-like method described in Ref. [52] to control the step-size when converging towards the saddle-point.

We also explore the effect of the initial deformation of the system on the overall performance of the saddle-search. This is known to affect the convergence and sampling efficiency of this class of methods [49, 50]. In the case of the dimer method, we try randomly deforming the system both locally, constraining the initial deformation within 2 lattice-parameters of the point defects, and globally, deforming the entire Active Volume. In ARTn, we controlled the maximum number of steps where the system is pushed in a direction determined at random at the start of the search. If the maximum number of steps is reached or if the softest eigenmode becomes imaginary (negative eigenvalue), then ARTn starts following this eigendirection.

We use the Ackland (2004)[54] semi-empirical potential for Fe. In the case of FeCr, we use a two-band potential where Fe interactions are based on Ackland (2004), as developed by Olsson [55] . Saddles are converged within a precision of 0.0285 eV/Å (0.01 eV per lattice parameter). All benchmarks use an OpenMP parallelization over 12 cores on 2x X5650 Intel Xeon CPUs (at 2.67 GHz).

First, we compare the performance of both frameworks in the case of a mono-vacancy in a 10-lattice-parameter Fe box (1999 atoms). Diffusion in such a system is characterized by a two-step process, which includes a stable state and a metastable state [56]. We launch 120 saddle searches from both states and a use a 2.7 lattice-parameter radius spherical active volume. In order to better understand the saddle-search process, we also ran modified versions of the algorithms on this system. For instance, we ran a version of FIRE where a Lanczos evaluation of the softest mode is performed at every step and where

the force is inverted in that eigendirection. We also ran a variant of the dimer method that minimizes $\mathbf{F} = \mathbf{F}_T - 2\mathbf{F}_p$ both when the curvature is concave and convex, where $\mathbf{F}_T$ is the total force on the dimer and $\mathbf{F}_p$ is the force parallel to the dimer. In the standard dimer algorithm, this quantity is minimized only when the curvature is convex.

Second, we compare the performance of both frameworks in the case of a (110) dumbbell SIA in a 10-lattice-parameter Fe box (2001 atoms). We launch 120 saddle searches and use a 4.0 lattice-parameter radius spherical active volume.

Third, we compare the performance of both frameworks in the case of 50 vacancies in a 10-lattice-parameter box (1950 atoms). We launch 120 saddle searches from one of the states generated by a SEAKMC run, using a 2.7 lattice-parameter radius spherical active volume.

Fourth, we compare the performance of both frameworks in the case of two 37-atom interstitial loops in a 20-lattice-parameter Fe box (16074 atoms). We launch 120 saddle searches from one of the states generated by a SEAKMC run, using a 4.5 lattice-parameter radius spherical active volume.

Fifth, we compare the performance of both frameworks in the case of two 37-atoms interstitial loops in a 20-lattice-parameter FeCr (10% Cr) box (16074 atoms). We launch 120 saddle searches from one of the states generated by a SEAKMC run. We use a 4.5 lattice-parameter radius spherical active volume.

2.3. Set-Up of on-the-fly KMC comparison

2.3.1. 50-vacancy Problem

We use the same 50-vacancy system as for the saddle-search comparison. We simulate the kinetics of aggregation at 323K, with a constant pre-factor of $5 \times 10^{12} s^{-1}$. We run simulations until the potential energy of the system reaches -7760 eV, which roughly corresponds to the energy when half the mono-vacancies have aggregated.

Previous comparisons of k-ART and SEAKMC looked at simulations of the aggregation of 50 Fe vacancies [5, 10]. However, while many simulation parameters were set-up as to make them as comparable as possible, some details were

overlooked. For instance, different randomly generated initial configurations were used in each k-ART run, while the SEAKMC runs all started from the same initial state. In this study, all initial states are identical. Furthermore, the number of saddle-searches performed were not tracked in the original studies. As our results show, this is an important parameter for the 50-vacancy problem.

2.3.2. Fe Intersitial Loops

In previous work, SEAKMC was used to discover how two interstitial loops with $< 111 >$ and $< -1 - 11 >$ Burgers vectors aggregate and transform to form either $< 111 >$, $< -1 - 11 >$ or $< 100 >$ loops [8]. Here, we perform these calculations using k-ART. We look at 37-intersitial loops in a 20-lattice parameter cubic box (for a total of 16074 atoms). The simulations are run at 573K, using the Ackland (2004) potential. In order to escape basins dominated by flickers and complete the transformation, we combined the bacMRM (with a 0.8 eV threshold to escape a super-basin) to a threshold on total RMSD displacement of each event (of 2.1 Å). Events with smaller displacements were ignored. This procedure ressembles that used to in Refs. [1, 26] to study SiC and sputtering of Au surfaces with variations of k-ART. We start all simulations in a configuration where the two loops are interlocked (see Fig. 8), which corresponds to a potential energy of -64404.05 eV.

3. Results and Discussion

3.1. Comparison of saddle-search methods

3.1.1. Fe Mono-vacancy

The results are shown in Tables 2 and 3. We see that the methods both find the same activation barriers for diffusion, i.e. 0.64 eV from the stable state and 0.09 eV from the metastable state, in agreement with previous studies. The dimer method also finds a second saddle-point at 1.07 eV and other higher-energy saddle-points. This corresponds to an event where an atom moves to a position half-way between two empty lattice sites. For both states, stable and metastable, the fastest method is the local dimer, closely followed by the ARTn implementation that limits the number of steps in the initial direction to 5. The

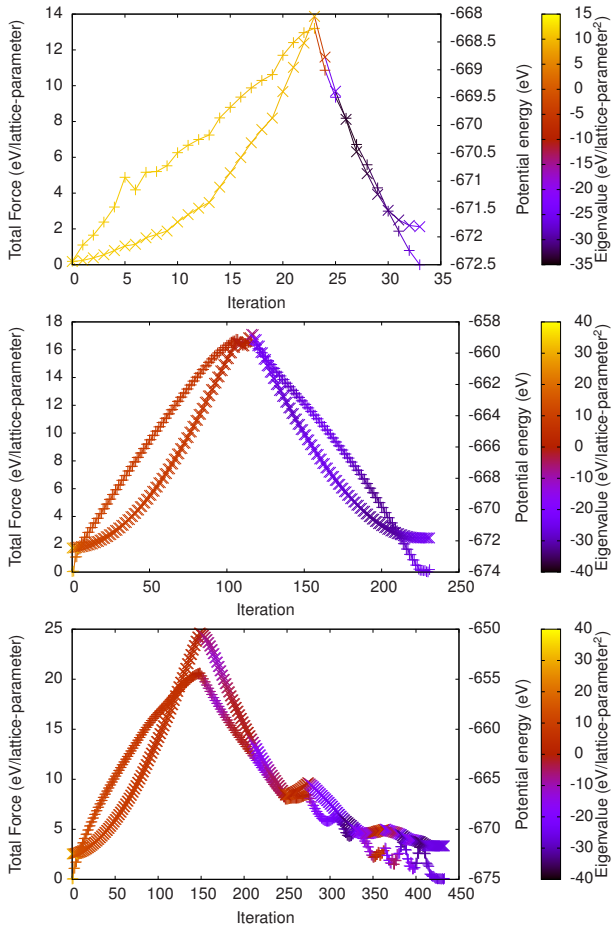

Figure 2: Total force (+) and potential energy (x) in typical saddle-searches for the Fe mono-vacancy system. The eigenvalue at each iteration is color-coded. The top figure corresponds to an ARTn search and the two bottom figures correspond to dimer searches. The the bottom figure differs because of a different random seed when initializing the saddle-search.

global dimer search is slower than the local version. The ARTn searches with a maximum of 20 steps pushing in the initial direction is the slowest of the search methods.

We show examples of saddle-search profiles with ARTn and the dimer method in the Fe vacancy case in Fig. 2. These profiles are typical of the saddle-searches for all the systems reported in this paper. The dimer and ARTn profiles are very similar. We see that the system's potential energy and forces increase until the eigenvalue of the softest mode becomes negative. Thereafter, both the potential energy and the forces decrease. This behavior does not seem to be dependent on the exact implementation of the saddle-search. Indeed, a FIRE+Lanczos-based implementation and a variant to the dimer method (see section 2.2) showed profiles very similar to those of ARTn and the standard conjugate-gradient-based dimer method. However, these two modified methods, that rely more heavily on relaxation in the perpendicular directions to the softest mode, often lead to searches that are trapped near the initial minimum, failing to converge to the saddle point.

One should note that a typical ART iteration is about 10 times more computationally expensive than a typical dimer iteration, because of the successive application of the Lanczos algorithm and tens of FIRE steps, compared to a few conjugate-gradient-based rotation iterations and translation attempts in a typical dimer-step. On the other hand, about 10 times less ARTn steps are needed to converge to the saddle.

We also notice that the time to converge to the saddle is related to the potential energy when a convexity is found. The higher the energy, the more iterations are necessary to reach the saddle-point. We also observe in the bottom panel of Fig. 2 that convergence is more difficult if, during the minimization phase of the saddle-search, a concavity is hit. These correspond to the red-colored points around iteration number 175, number 250 and number 350. Nevertheless, by following the softest mode, the dimer method manages to find the saddle-point. However, one should keep in mind that hitting a concavity might indicate that we are no longer attracted to the saddle we were initially pursuing. This might

lead to saddles that are not connected to the original minimum.

3.1.2. Fe SIA

The results are shown in Table 4. We see that both methods find the 0.31 eV barrier for diffusion and the 0.41 eV barrier to rotate to a $1/2 < 111 >$ configuration, in agreement with the literature. In this implementation, the dimer method samples more high-energy saddle-points. This is particularly true if the initial deformation is across the whole active volume. We have also tested this benchmark using a smaller step-size (70% smaller, not shown in the table). While the execution was much slower (it was nearly four times slower), all searches resulted in 0.31 eV and 0.41 eV barriers.

Both implementations of ARTn perform better than the local dimer implementation (48 and 51 seconds vs 54 seconds, respectively). These three local deformation algorithms perform much better than the global deformation dimer, that is nearly 8 times slower and has a 11% lower success rate than the local methods.

The search profiles are similar to those of the mono-vacancy, plotted in Fig. 2. We see that searches that converge the slowest or that do not lead to a saddle are those that reach very high levels of potential energy and forces, which happens most often when globally deforming the active volume. We believe that the use of a bigger active volume for the SIA than the mono-vacancy (4 lattice-parameters vs 2.7) explains why a global deformation is so inefficient.

3.1.3. 50 Fe Vacancies

The results are shown in Table 5. The fastest method is the local dimer (126 seconds), followed by the ARTn implementation that limits the number of steps in the initial direction to 5 (156 seconds). The global dimer search is slower than the local version (166 seconds). The ARTn searches with a maximum of 20 steps pushing in the initial direction is the slowest of the search methods (260 seconds).

The distribution of saddles for each method is shown in Figure 3. They are very similar, no matter which method is selected. The local dimer method

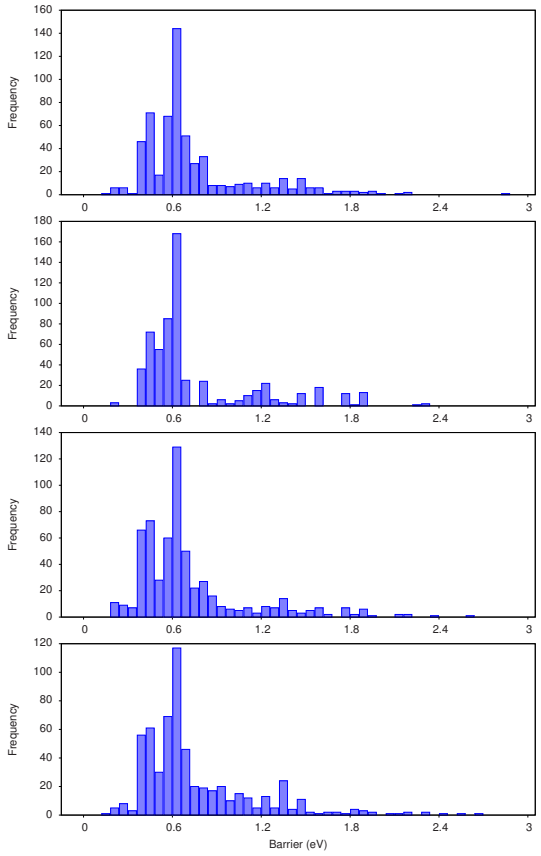

Figure 3: Distribution of saddles found by the dimer method (top two panels) and ARTn (two bottom panels) for an Fe 50-vacancy configuration. From top to bottom: global dimer, local dimer, ARTn limited to 5 initial displacement steps, ARTn limited to 20 initial displacement steps.

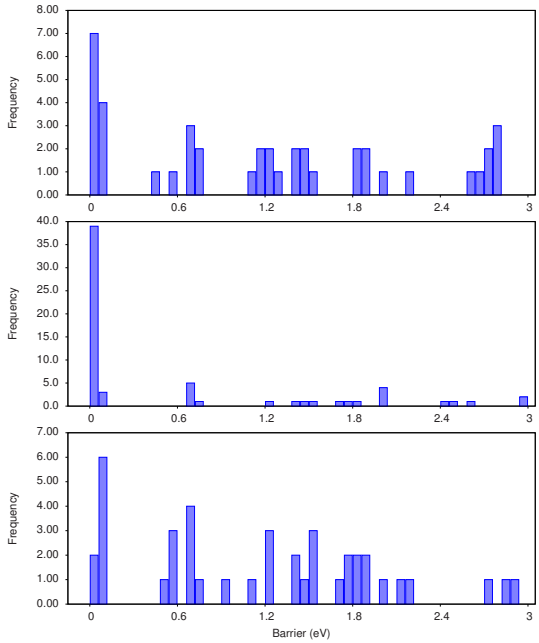

Figure 4: Distribution of saddles found by the dimer method (top panel) and ARTn (two bottom panels) for the Fe interstitial loops system. From top to bottom: global dimer, ARTn limited to 3 initial displacement steps, ARTn limited to 20 initial displacement steps.

generates the most narrow distribution (with a peak frequency of 168) and the ARTn runs limited to 20 initial displacements steps have the broadest distribution (with a peak frequency of 117). Thus, while the local dimer generates saddle points more rapidly, its sampling is not as broad as that of other alternatives. As in the SIA case, the search profiles are similar to those of the mono-vacancy, plotted in Fig. 2.

We notice that the relative efficiency of the saddle-search implementations in this system is similar to the mono-vacancy system. Even though they interact, the energy landscape formed by the 50 vacancies has similarities with the landscape of a single vacancy.

3.1.4. Fe interstitial loops

The results are shown in Table 6. We did not implement a local version of the dimer method. The fastest method is the ARTn implementation that

limits the number of steps in the initial direction to 3 (258 seconds). The global version of the dimer does better than the ARTn implementation that limits the number of steps in the initial direction to 20 (339 seconds vs 349 seconds).

In Fig. 4, we show the distribution of saddles found by the three methods. Overall, the global dimer version and the ARTn runs that limit initial displacements steps to 20 are very similar. They each show a double peak at low barriers with 2-7 events (between 0 and 0.06 eV) and many peaks of size 1 to 4 saddles at higher energies. When limiting ARTn searches to 3 steps in the initial direction, we observe a much larger peak at low-energy (between 0 and 0.03 eV), with 39 events, and a few occurrences at high energies. This latter implementation seems to favor low-energy barriers.

In this case, localizing the initial deformation does not seem to greatly affect the efficiency of the search. One should note that both the ARTn and dimer searches lead to states where the whole volume has displaced atoms at the saddle-point, which may explain why they perform similarly.

3.1.5. FeCr interstitial loops

The results are shown in Table 7. We did not implement a local version of the dimer method. The fastest method is the ARTn implementation that limits the number of steps in the initial direction to 20 (381 seconds), closely followed by the ARTn implementation that limits the number of steps in the initial direction to 3 (383 seconds) . The global version of the dimer does much worse than ARTn (2248 seconds). The dimer also has a much smaller success rate (5.6%) than ARTn (about 30%).

In Fig. 5, we show the distribution of saddles found by three methods. The dimer method only found 8 events, scattered across the whole spectrum. The two ARTn implementation have somewhat similar spectra. When limiting ARTn searches to 3 steps in the initial direction, we observe more low-energy events (from 0 to 0.6 eV). This latter implementation seems to favor low-energy barriers, as in the pure Fe case.

The addition of Cr to the system seems to create problems when activating

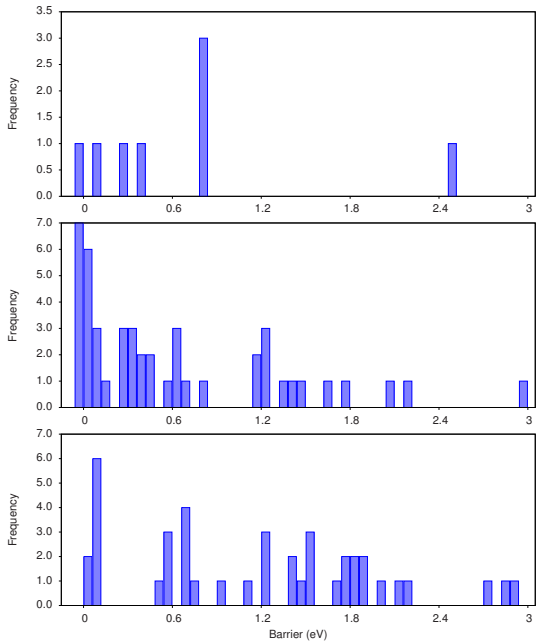

Figure 5: Distribution of saddles found by the dimer method (top panel) and ARTn (two bottom panels) for the FeCr interstitial loops system. From top to bottom: global dimer, ARTn limited to 3 initial displacement steps, ARTn limited to 20 initial displacement steps.

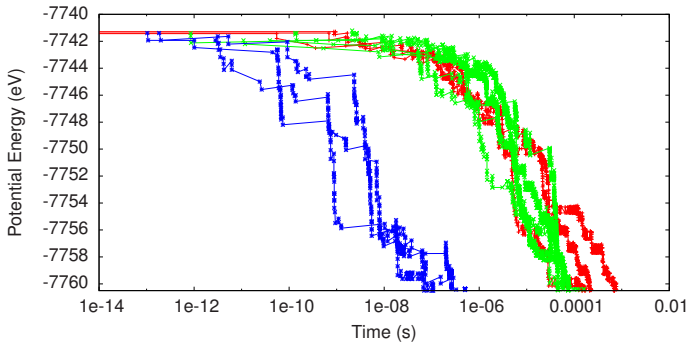

Figure 6: Potential energy of 50 mono-vacancies aggregating at 323K. Eight of these curves have a large overlap.The three red curves correspond to k-ART runs and the five green curves correspond to SEAKMC runs that sample 2 events per defect per step. The three blue curves on the left that do not overlap with the others are SEAKMC runs that sample 24 events per defect per step

the whole volume. This might be related to the presence of optical phonons, which are absente in the pure Fe, and to the many-body nature of the two-band model, which can add roughness to the second derivatives of the potential energy surface. The effect of these vibrational modes and sharp changes in forces on the eigenmode-following search are probably enhanced when deforming the lattice rapidly over a large radius.

3.2. On-the-fly KMC comparison

3.2.1. Fe 50-vacancy problem

We tracked the generation of saddle points in k-ART and SEAKMC when simulating the aggregation of 50 vacancies in a 1950-atom Fe box. A notable feature of SEAKMC is that it has an overall success rate of saddle-convergence of 95%. This is possible because searches are centered directly on the crystallographic defects. In comparison, k-ART has a success rate of 35% for generic-event generation (before analysis to exclude multiple identical events) and a 90% specific-event generation success rate.This is because the saddle-searches on not centered on crystallographic defects, but rather on atoms with a connectivity graph different from that of an atom in a perfect crystal. As a consequence, while there are 50 defects(vacancies) in the system, the NAUTY-based topol-

ogy analysis identifies, on average, 250 different topologies in the system, which must all be explored by k-ART. Simple geometrical considerations imply that most topologies represent atoms that sit far from the defects, near the limit of the region where atoms are included in the topology-analysis, i.e. a sphere with a radius of 5.6 Å. Centering searches on these atoms is not likely to generate events, especially if the displacements are not limited to an Active Volume. As for failures of specific-event generation, this is related to failures to execute the correct geometrical transformation from the generic event's coordinates to a specific atomic site, which relies on the graph canonical representation, which can be problematic for configurations where inversion symmetries play a role.

In such an off-equilibrium system, the kinetics are largely determined by the presence of high-energy barriers that separate the system from relaxation events (i.e. the replenish-and-relax model). If one misses a low-lying energy barrier that can eventually lead to relaxation and instead chooses to execute a high-energy barrier, the kinetics will be unphysically slowed down. Good sampling of the potential energy landscape between relaxation steps (i.e. the aggregation of a vacancy with a cluster) is crucial for good kinetics.

Considering that 20 generic searches per new topology were launched in k-ART, that the total number of topologies visited during the runs is about 8600, that the number of KMC steps for each run is roughly 700 and a 35% success rate, we estimate that, on average, approximately 85 events were generated every step. Taking into consideration the new events created by the refinement of specific events, this is comparable to the 100 event searches (2 searches per defect) launched in SEAKMC.

In Fig. 6 , we show five SEAKMC runs with 2 searches per defect and three k-ART runs, that are stopped after the system reaches a potential energy of -7760 eV, which roughly corresponds to the energy when half the mono-vacancy remain. All the simulations are in good agreement. We note that some SEAKMC runs permitted the execution of events with negative energy barriers. Although these are unphysical within TST, they usually correspond to instabilities in the potential energy. The good agreement of the runs in

Fig. 6 indicate that executing these events does not play an important role in determining the kinetics.

Furthermore, we tried two different activation volumes (2.7 Å and 4.0Å). This does not seem to affect the aggregation time. However, we noticed that the two runs with an activation volume of 2.7 Å took about 1600 steps to reach -7760 eV, while those with a larger active volume took about 300 steps. This is probably due to the fact that we used a global dimer search. Indeed, when initial deformations are global, the saddle search involves configurations of higher energy, as we explain in Sec. 3.1.1. The larger the active volume, the greater the effect. This favors finding saddles with higher energies, which, if using a limited sampling of events, can help escape flickers. Indeed, we tried such a simulation using the ARTn, which performs localized initial deformations, within SEAKMC (not shown). We limited the number of steps in the initial direction to 5, as to favor low-energy barriers as much as possible, and found that the system oscillates between states separated by low-barriers for a long time.

It thus appears that using an active volume and a global dimer method indirectly permits handling flickers by favoring events with larger barriers and displacements. Although, unlike the bacMRM, this is not a rigorous method, the results in Fig. 6 seem to indicate that it produces the right kinetics. It also indicates that the two different sets of approximations implemented in k-ART and SEAKMC, described in Table 1, produce very similar kinetics, at least for this system. We note that the use of a too aggressive basin threshold leads to slower relaxation (not shown in the figures), because it leads to less sampling and thus can miss low-barrier relaxation events.

From the standpoint of computational efficiency, the SEAKMC took roughly 70 cpu-hours to reach -7760 eV, while the k-ART runs took 2500. This is mostly due to the use of the Active Volume, which permits force calculations across a much smaller number of atoms (165 vs 1950). The rest of the acceleration is associated with the larger success rates of saddle-searches. We note that for a system of such size, implementing localized forces calculations that dynami-

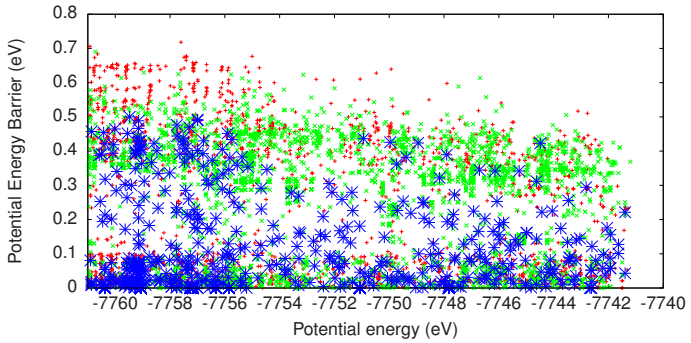

Figure 7: Potential barriers energy of executed events of 50 mono-vacancies aggregating at 323K as a function of the potential energy of the local minimum. The red (+) points are k-ART runs, the green (X) points are SEAKMC runs sampling 2 events per defect per step and large blue stars are SEAKMC runs sampling 24 events per defect per step.

cally include new participants, as described in Refs. [2, 3], does not produce a significant speedup, because of the computational overhead of such a tactic.

We also ran a group of three simulations that sample 1200 events per step (24 events per defect), twelve times more than the other SEAKMC and k-ART simulations. Because of the large cost of such sampling, we only performed them with SEAKMC. The results show that the kinetics are 1000 times faster. Even if we assume that adding 12 times more events to the catalog (thus increasing the rates by a similar proportion) amounts to oversampling, the kinetics are nearly 100 times faster at this level of sampling. These faster runs correspond to the aggregation time observed in Ref. [5]. This reconciles the apparent mismatch between previous publications.

Fig. 7 explains this behavior. It shows the barriers that are executed as a function of the potential energy of the local minimum. Since the potential energy decreases nearly-monotonously with time, reading this graph from right to left gives a measure of the activation barriers that characterize relaxation. We see that the k-ART and SEAKMC runs with comparable sampling (100 events per step) execute events with similar barriers. The SEAKMC runs that sample 1200 events per step are characterized by smaller barriers (roughly 0.25 eV lower) for

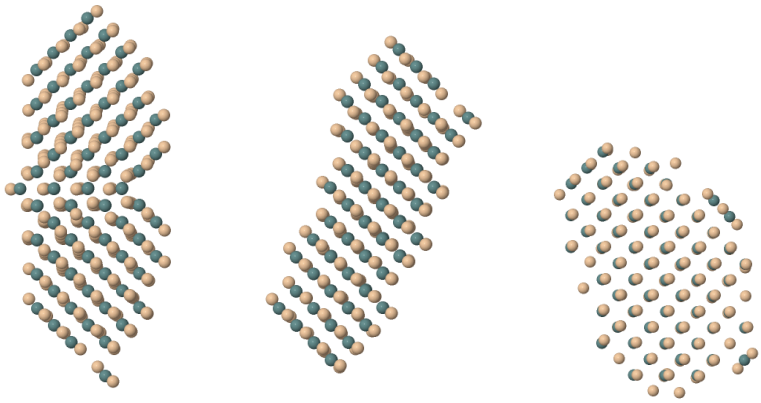

Figure 8: Illustration of the initial state for the Fe interstitial-loop runs (upper left) and two outcomes: transformation to a $< 111 >$ loop (upper right) and to a $< 100 >$ loop (bottom). All illustrations are from the $< 100 >$ perspective. Beige spheres are interstitials and blue spheres are vacancies.

all configurations. In other words, by sampling more, the algorithm found events with lower barriers that nonetheless permit large displacements without many flickers. This is an indication that we were in fact under-sampling the system in the other runs.

3.2.2. Fe Interstital-loop problem

We show illustrations of the initial configuration for our runs (two $1/2 < 111 >$ loops) and of two outcomes after a transformation (a $1/2 < 111 >$ loop and a $< 100 >$ loop) in Fig. 8. In order to complete the transformation in a shorter amount of states, as explained in the methodology, all events with an RMSD smaller than 2.1 Å are ignored. If events with lower displacements, e.g. 1.0 Å, are included in k-ART, flickering motions are created, leading to super-basins with several hundred states before a transformation can be executed. By imposing the threshold on displacements, transformations were executed in k-ART within 140 to 190 KMC steps, not including flickers treated by the bacMRM. We note that this is larger than the number of steps in the SEAKMC runs described in Ref. [8], that took less than 100 KMC steps to execute. We

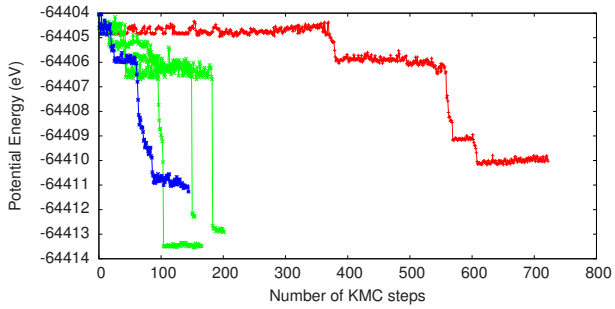

Figure 9: The potential energy of Fe loop system as a function of the number of KMC steps. The red line (+) corresponds to a k-ART run with a RMSD displacement threshold of 1.0Å on events (see text) that leads to a $< 100 >$ loop. The other runs have a 2.0Å displacement threshold. The green (X) lines lead to a $1/2 < 111 >$ loop, while the blue (stars) leads to a $< 100 >$ loop.

believe that this is related to the use of global initial deformations, which favors events with larger barriers and displacements. We also note that sampling a large number of events per step (more than 120) also leads to flickering in SEAKMC. In other words, SEAKMC handles the flicker problem by imposing large initial deformations in the saddle search and by limiting the sampling of the system at every step.

Both k-ART and SEAKMC predict transformations to $< 100 >$ and $1/2 < 111 >$. We show the potential-energy evolution during these runs in Fig.9. Over these five runs, the highest barrier crossed was 0.81 ± 0.23 eV. In Ref. [8], averaging over more than one hundred SEAKMC simulations, the authors found 0.73 ± 0.11 eV. Furthermore, the atomistic mechanism that these two algorithms describe were confirmed through MD simulations. Overall both methods come to similar conclusions, even though they each use a very different set of approximations.

The k-ART simulations took 2000-2500 cpu-hours to execute the transformation, while SEAKMC took about 71 cpu-hours. This ratio is in fact similar to that found for the 50-vacancy case. This is mostly due to the use of the Active Volume, that permits force calculations to be limited to 1050 atoms, compared to 16074 atoms in the case of k-ART. We note that we were able to

implement a local force calculation algorithm for this potential within k-ART (as described in Refs. [2, 3]), but the overall speed was comparable to the use of the optimized force computation library of IMD, as implemented in Ref.[2].

4. General Remarks

4.1. Saddle-Search Methods

In order to converge quickly to the saddle, one should attempt to limit the potential energy and forces of the system when it reaches the convexity of the potential energy surface. A straightforward way to do this is to limit the number of atoms which are involved in the initial deformation of the system. Indeed, in our benchmarks, localized initial deformations have led to faster convergence to the saddle-point than deformations across the whole active volume. Also, limiting the number of steps in the initial random displacement with ARTn can accelerate convergence in some systems, such as the mono-vacancy, the 50-vacancy and the pure Fe interstitial-loop systems presented here. However, this has not helped increased convergence in the SIA case, nor the FeCr interstitial-loop case.

This general assessment, that locally constraining initial displacements generally leads to faster saddle-searches and good sampling of exit states, was noted and quantified in previous studies [49, 50], albeit it was not shown to apply both to the dimer method and ARTn, nor specifically to Fe and FeCr defects.

The increase and subsequent relaxation of the potential energy indicate that these methods do not follow the minimum energy path (MEP) to the saddle-point. One might think that this is related to an insufficient relaxation in the directions perpendicular to the softest eigenmode. However, the same behavior was observed when modifying the algorithms to emphasize perpendicular relaxation. Why is this? The literature provides some insights to this effect. Take, for instance, Refs. [34, 57, 58]. The first is the original dimer paper and the other two describe the eigenvector-following method and the Gentlest-Ascent Dynamics (GAD), which are closely related to the dimer and ARTn methods. These studies present tests on simple, 2-dimensional potentials. In all case,

if the initial deformation is not sufficiently parallel to the MEP, the potential energy will follow an increase and subsequent decrease when searching for a saddle-point. In many-dimensions, is it virtually impossible to generate a random displacement that has a large parallelity with the MEP, thus explaining why we observe profiles similar to those illustrated in Fig. 2.

We note that Ref. [58] also provides an alternative, open-ended, method to find saddle-points that closely follow the MEP, GAD molecular dynamics (GAD-MD). In this case, one generates an approximation of the eigendirection of lowest curvature using the variational principle and runs a molecular dynamics simulation, applying $\mathbf{F} = \mathbf{F}_T - 2\mathbf{F}_p$, as in the dimer method and the GAD. At every time step, the approximation of the eigendirection is updated using the Euler method. This special bias of the potential destabilizes the energy landscape's minima and stabilizes the saddle-points, while favoring low-lying free-energy-states. In their test on a 2-dimensional surface, the trajectories seem to stay close to the MEP. We also note that in Ref. [34], the authors suggest that dimer searches should start from MD configurations in order to make the trajectories as physically relevant as possible. In that same spirit, we believe that executing GAD-MD until a convexity is discovered and then starting an eigenmode-following method might permit convergence to saddle points relatively efficiently while staying near the MEP. However, such a method may favor low-energy saddles even more than do traditional methods, which may make sampling less efficient. While such an implementation is out of the scope of this work, such development would be of interest, at least from a pedagogical perspective, even if not from the perspective of kinetics simulation.

4.2. On-the-fly KMC

It seems that both sets of approximations implemented in k-ART and SEAKMC lead to similar kinetics. We showed that SEAKMC treatment of flickers implicitly depends on the use of the global dimer method and somewhat undersampling the system. While this has worked well for the two systems described in this study, it is quite possible it might not work in other cases. For example,

our calculations in FeCr indicate that using the global dimer method is not adequate. Preliminary simulations of interstitial loops in FeCr show that using such a technique widely overestimates the activation barriers.

It is preferable, in our opinion, to use saddle-search methods based on local deformations (or MD) and afterwards discard events which will probably not be important to the kinetics (such as events with very small symmetric barriers and events with small RMSD displacements). This should increase the chance of finding low-barrier events that lead to significant changes in the system. This type of procedure seems widespread in the literature. Another approach is to massively increase the temperature in the hope that higher-energy barriers will be selected and exit the flickering super-basin [1].

However, discarding events in order to minimize flickering is a procedure that can significantly distort kinetics. Two methods exist in kinetic-ART to solve this problem: TABU and the bacMRM. In Ref. [59], we show that these methods generate very similar kinetics. If possible, one should use the bacMRM, since it is a more rigorous approach. If super-basins are very large, this might not be the case, and the use of TABU may be preferable to discarding low-displacement events.

We also note that many of the algorithms are not fundamentally tied to all the approximations described in Table 1. For example, detailed balance and checks for duplicity of events can be implemented in SEAKMC through bookkeeping. Similarly, the performance gap between k-ART and SEAKMC could be narrowed by introducing the active volume approximation in k-ART. We should note that in Si, k-ART's local force implementation led to essentially no cost increase with system-size [3] above 5000-atoms. In such systems, the active volume advantage would probably be much smaller.

Another strong feature of k-ART is its proven ability to simulate the kinetics of amorphous materials. It is not clear if the Active Volume approximation is appropriate for these highly disordered materials. In addition, while the topology-based criteria that coordinates saddle-searches and cataloging is completely independent of crystallinity, *ad hoc* procedures for coordinating saddle-

searches would probably need to be implemented in SEAKMC. Further work is needed to clarify this open issue.

5. Conclusion

In this work, we have have compared and contrasted two state-of-the-art aKMC methods, k-ART and SEAKMC. Firstly, we looked at the relative performance of the dimer method and ARTn, two open-ended saddle-search methods in simulations related to point defects, interacting vacancies and interstitial loops in Fe and FeCr. Secondly, we compared the kinetics of aggregation of 50 vacancies in a 1950-atom Fe box. Finally, we compared the kinetics of transformation of $1/2 < 111 >$ loops in Fe to either a $< 111 >$ loop or a $< 100 >$ loop.

Generally, the performance of the dimer method and ARTn are similar if one ensures that the initial displacements during a search are limited to a localized subset of atoms. This generally permits faster convergence, a better success rate, and increases the probability of finding low-barrier saddle-points, which are of crucial importance for kinetics.

k-ART and SEAKMC simulations of the 50-vacancy system are in agreement if one is careful to generate the same number of events per KMC step. We also discovered that under-sampling by a factor twelve leads to aggregation 1000 times slower, through the discovery of higher-barrier events that lead to relaxation. The two algorithms also draw the same picture for the kinetics of the dislocation loops.

This confirms that the different set of approximations that these algorithms implement were appropriate for this class of systems. It also indicates that both methods could be improved by adapting some elements from one another. For example, k-ART could significantly accelerate simulations by introducing active volumes, while the use of the bacMRM in SEAKMC would make it more rigorous when the energy landscape is dominated by flickers.

6. acknowledgements

LKB thanks Alexander V. Barashev for insightful discussions.

Research at the Oak Ridge National Laboratory sponsored by the U.S. Department of Energy, Office of Basic Energy Sciences, Materials Sciences and Engineering Division, "Center for Defect Physics," an Energy Frontier Research Center. LKB acknowledges a fellowship awarded by the Fonds Québécois de recherche Nature et Technologies.

k-ART is available upon request to Normand Mousseau and SEAKMC is available upon request to the authors.

References

[1] H. Jiang, C. Jiang, D. Morgan, I. Szlufarska, Accelerated atomistic simulation study on the stability and mobility of carbon tri-interstitial cluster in cubic sic, Computational Materials Science 89 (2014) 182–188.

[2] L. K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, N. Mousseau, Kinetic activation-relaxation technique, Physical Review E 84 (4) (2011) 046704.

[3] J.-F. Joly, L. K. Béland, P. Brommer, F. El-Mellouhi, N. Mousseau, Optimization of the kinetic activation-relaxation technique, an off-lattice and self-learning kinetic monte-carlo method, Journal of Physics: Conference Series 341 (1) (2012) 012007.

[4] P. Brommer, N. Mousseau, Comment on mechanism of void nucleation and growth in bcc fe: Atomistic simulations at experimental time scales, Physical Review Letters 108 (21) (2012) 219601.

[5] H. Xu, R. E. Stoller, Y. N. Osetsky, Cascade defect evolution processes: Comparison of atomistic methods, Journal of Nuclear Materials 443 (1) (2013) 66–70.

[6] H. Xu, Y. N. Osetsky, R. E. Stoller, Simulating complex atomistic processes: On-the-fly kinetic monte carlo scheme with selective active volumes, Physical Review B 84 (13) (2011) 132103.

[7] H. Xu, Y. N. Osetsky, R. E. Stoller, Self-evolving atomistic kinetic monte carlo: fundamentals and applications, Journal of Physics: Condensed Matter 24 (37) (2012) 375402.

[8] H. Xu, R. E. Stoller, Y. N. Osetsky, D. Terentyev, Solving the puzzle of¡ 100¿ interstitial loop formation in bcc iron, Physical review letters 110 (26) (2013) 265503.

[9] S. T. Chill, G. Henkelman, Molecular dynamics saddle search adaptive kinetic monte carlo, The Journal of Chemical Physics 140 (21) (2014) 214110.

[10] P. Brommer, L. K. Béland, J.-F. Joly, N. Mousseau, Understanding long-time vacancy aggregation in iron: a kinetic activation-relaxation technique study (2014). arXiv:arXiv:1407.7694.

[11] A. Kara, O. Trushin, H. Yildirim, T. S. Rahman, Off-lattice self-learning kinetic monte carlo: application to 2d cluster diffusion on the fcc (111) surface, Journal of Physics: Condensed Matter 21 (8) (2009) 084213.

[12] G. Nandipati, Y. Shim, J. G. Amar, A. Karim, A. Kara, T. S. Rahman, O. Trushin, Parallel kinetic monte carlo simulations of ag (111) island coarsening using a large database, Journal of Physics: Condensed Matter 21 (8) (2009) 084214.

[13] A. Latz, L. Brendel, D. E. Wolf, A three-dimensional self-learning kinetic monte carlo model: application to ag (111), Journal of Physics: Condensed Matter 24 (48) (2012) 485005.

[14] H. Jónsson, Simulation of surface processes, Proceedings of the National Academy of Sciences 108 (3) (2011) 944–949.

[15] G. Henkelman, H. Jónsson, Long time scale kinetic monte carlo simulations without lattice approximation and predefined event table, The Journal of Chemical Physics 115 (21) (2001) 9657–9666.

[16] O. Trushin, A. Karim, A. Kara, T. S. Rahman, Self-learning kinetic monte carlo method: application to cu (111), Physical Review B 72 (11) (2005) 115401.

[17] A. Pedersen, G. Henkelman, J. Schiøtz, H. Jónsson, Long time scale simulation of a grain boundary in copper, New Journal of Physics 11 (7) (2009) 073034.

[18] L. Xu, G. Henkelman, Adaptive kinetic monte carlo for first-principles accelerated dynamics, The Journal of chemical physics 129 (11) (2008) 114104.

[19] L. Xu, D. Mei, G. Henkelman, Adaptive kinetic monte carlo simulation of methanol decomposition on cu (100), The Journal of chemical physics 131 (24) (2009) 244520.

[20] L. Karssemeijer, A. Pedersen, H. Jónsson, H. Cuppen, Long-timescale simulations of diffusion in molecular solids, Physical Chemistry Chemical Physics 14 (31) (2012) 10844–10852.

[21] E. A. Bleda, X. Gao, M. S. Daw, Calculations of diffusion in fcc binary alloys using on-the-fly kinetic monte carlo, Computational Materials Science 43 (4) (2008) 608–615.

[22] F. El-Mellouhi, N. Mousseau, L. J. Lewis, Kinetic activation-relaxation technique: An off-lattice self-learning kinetic monte carlo algorithm, Physical Review B 78 (15) (2008) 153202.

[23] N. Mousseau, L. K. Béland, P. Brommer, J.-F. Joly, F. El-Mellouhi, E. Machado-Charry, M.-C. Marinica, P. Pochet, The activation-relaxation technique: Art nouveau and kinetic art, Journal of Atomic, Molecular, and Optical Physics 2012.

[24] J.-F. Joly, L. K. Béland, P. Brommer, N. Mousseau, Contribution of vacancies to relaxation in amorphous materials: A kinetic activation-relaxation technique study, Physical Review B 87 (14) (2013) 144204.

[25] C. Scott, S. Blackwell, L. Vernon, S. Kenny, M. Walls, R. Smith, Atomistic surface erosion and thin film growth modelled over realistic time scales, The Journal of chemical physics 135 (17) (2011) 174706.

[26] C. Scott, R. Smith, Sputtering of au surfaces at realistic dose rates using molecular dynamics and on-the-fly kinetic monte carlo, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 303 (2013) 192–195.

[27] L. K. Béland, Y. Anahory, D. Smeets, M. Guihard, P. Brommer, J.-F. Joly, J.-C. Pothier, L. J. Lewis, N. Mousseau, F. Schiettekatte, Replenish and relax: Explaining logarithmic annealing in ion-implanted c-si, Physical Review Letters 111 (10) (2013) 105502.

[28] L. K. Béland, N. Mousseau, Long-time relaxation of ion-bombarded silicon studied with the kinetic activation-relaxation technique: Microscopic description of slow aging in a disordered system, Physical Review B 88 (21) (2013) 214201.

[29] S. Glasstone, K. J. Laidler, H. Eyring, The theory of rate processes: the kinetics of chemical reactions, viscosity, diffusion and electrochemical phenomena, McGraw-Hill Book Company, Incorporated, 1941.

[30] G. H. Vineyard, Frequency factors and isotope effects in solid state rate processes, Journal of Physics and Chemistry of Solids 3 (1) (1957) 121–127.

[31] A. F. Voter, F. Montalenti, T. C. Germann, Extending the time scale in atomistic simulation of materials, Annual Review of Materials Research 32 (1) (2002) 321–346.

[32] G. Henkelman, B. P. Uberuaga, H. Jónsson, A climbing image nudged elastic band method for finding saddle points and minimum energy paths, The Journal of Chemical Physics 113 (22) (2000) 9901–9904.

[33] M. R. So, A. F. Voter, et al., Temperature-accelerated dynamics for simulation of infrequent events, The Journal of Chemical Physics 112 (21) (2000) 9599–9606.

[34] G. Henkelman, H. Jónsson, A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives, The Journal of chemical physics 111 (15) (1999) 7010–7022.

[35] R. A. Horn, C. R. Johnson, Matrix analysis, Cambridge university press, 2012.

[36] L. J. Munro, D. J. Wales, Defect migration in crystalline silicon, Physical Review B 59 (6) (1999) 3969.

[37] G. Barkema, N. Mousseau, Event-based relaxation of continuous disordered systems, Physical review letters 77 (21) (1996) 4358.

[38] R. Malek, N. Mousseau, Dynamics of lennard-jones clusters: A characterization of the activation-relaxation technique, Physical Review E 62 (6) (2000) 7723.

[39] R. Olsen, G. Kroes, G. Henkelman, A. Arnaldsson, H. Jónsson, Comparison of methods for finding saddle points without knowledge of the final states, The Journal of chemical physics 121 (20) (2004) 9776–9792.

[40] A. Heyden, A. T. Bell, F. J. Keil, Efficient methods for finding transition states in chemical reactions: Comparison of improved dimer method and partitioned rational function optimization method, The Journal of chemical physics 123 (22) (2005) 224101.

[41] E. Machado-Charry, L. K. Béland, D. Caliste, L. Genovese, T. Deutsch, N. Mousseau, P. Pochet, Optimized energy landscape exploration using the ab initio based activation-relaxation technique, The Journal of chemical physics 135 (3) (2011) 034102.

[42] Y. Fan, A. Kushima, S. Yip, B. Yildiz, Mechanism of void nucleation and growth in bcc fe: atomistic simulations at experimental time scales, Physical review letters 106 (12) (2011) 125501.

[43] V. J. Bhute, A. Chatterjee, Building a kinetic monte carlo model with a chosen accuracy, The Journal of chemical physics 138 (24) (2013) 244112.

[44] V. J. Bhute, A. Chatterjee, Accuracy of a markov state model generated by searching for basin escape pathways, The Journal of chemical physics 138 (8) (2013) 084103.

[45] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, P. Gumbsch, Structural relaxation made simple, Physical review letters 97 (17) (2006) 170201.

[46] B. D. McKay, et al., Practical graph isomorphism, Department of Computer Science, Vanderbilt University, 1981.

[47] B. D. McKay, Nauty user?s guide (version 2.4), Computer Science Dept., Australian National University.

[48] B. Puchala, M. L. Falk, K. Garikipati, An energy basin finding algorithm for kinetic monte carlo acceleration, The Journal of chemical physics 132 (13) (2010) 134104.

[49] A. Pedersen, S. F. Hafstein, H. Jónsson, Efficient sampling of saddle points with the minimum-mode following method, SIAM Journal on Scientific Computing 33 (2) (2011) 633–652.

[50] A. Pedersen, M. Luiser, Bowl breakout: Escaping the positive region when searching for saddle points, The Journal of Chemical Physics 141 (2) (2014) 024109.

[51] M.-C. Marinica, F. Willaime, N. Mousseau, Energy landscape of small clusters of self-interstitial dumbbells in iron, Physical Review B 83 (9) (2011) 094119.

[52] E. Cances, F. Legoll, M.-C. Marinica, K. Minoukadeh, F. Willaime, Some improvements of the activation-relaxation technique method for finding transition pathways on potential energy surfaces, The Journal of chemical physics 130 (11) (2009) 114711.

[53] P. Császár, P. Pulay, Geometry optimization by direct inversion in the iterative subspace, Journal of Molecular Structure 114 (1984) 31–34.

[54] G. Ackland, M. Mendelev, D. Srolovitz, S. Han, A. Barashev, Development of an interatomic potential for phosphorus impurities in α-iron, Journal of Physics: Condensed Matter 16 (27) (2004) S2629.

[55] P. Olsson, J. Wallenius, C. Domain, K. Nordlund, L. Malerba, Two-band modeling of α-prime phase formation in fe-cr, Physical Review B 72 (21) (2005) 214119.

[56] L. Malerba, M. Marinica, N. Anento, C. Björkas, H. Nguyen, C. Domain, F. Djurabekova, P. Olsson, K. Nordlund, A. Serra, et al., Comparison of empirical interatomic potentials for iron applied to radiation damage studies, Journal of Nuclear Materials 406 (1) (2010) 19–38.

[57] J. C. Mauro, R. J. Loucks, J. Balakrishnan, A simplified eigenvector-following technique for locating transition points in an energy landscape, The Journal of Physical Chemistry A 109 (42) (2005) 9578–9583.

[58] A. Samanta, E. Weinan, Atomistic simulations of rare events using gentlest ascent dynamics, The Journal of chemical physics 136 (12) (2012) 124104.

[59] G. K. N'Tsouaglo, L. K. Béland, J.-F. Joly, P. Brommer, N. Mousseau, P. Pochet, Probing potential energy surface exploration strategies for complex systems (2014). `arXiv:arXiv:1407.7371`.

Table 1: The approximations that k-ART and SEAKMC use to accelerate on-the-fly off-lattice kMC. They correspond to the steps described by the flowchart of Fig. 1. There are no major approximations to the global minimization.

Identification of local configurations	
k-ART	SEAKMC
Local configurations are well described by topology of connectivity graph.	Local configs are well described by distances of atoms from the perfect lattice.

Association of events to local configurations	
k-ART	SEAKMC
Events can be catalogued using topology.	Only a subset of atoms are necessary to handle elastic effects on saddle- point.
Events can be mapped using graph canonical representation.	Sometimes neglect checks for duplicity of events.
Most important events are sampled when encountering a new topology.	Neglect detailed balance. (Formally, a large number of searches will still lead to micro-reversibility.)
Expansive, but incomplete, catalog is rarely resampled.	Catalog is resampled with uncorrelated random searches at every KMC step to compensate limited number of searches.
hTST	hTST
Constant prefactor	Constant prefactor

Event execution	
k-ART	SEAKMC
Poisson process.	Poisson process.
Uncorrelated processes.	Uncorrelated processes.
Mean-passage time correctly describes kinetics of flickers	

Table 2: Benchmarks for 120 saddle-searches from stable Fe mono-vacancy configuration.

Search method	Dimer	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local	Local
Max steps in init. direction	0	0	5	20
Success rate (%)	100	99.2	100	99.2
Execution time (s)	37	29	30	50
Saddles at 0.64 eV	104	111	120	119
Saddles at 1.07 eV	14	6	0	0
Other saddles	2	2	0	

Table 3: Benchmarks for 120 saddle-searches from metastable Fe mono-vacancy configuration.

Search method	Dimer	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local	Local
Max steps in init. direction	0	0	5	20
Success rate (%)	98.3	99.2	89.2	90
Execution time (s)	126	73	78	111
Saddles at 0.09 eV	118	114	107	108
Other saddles	0	5	0	0

Table 4: Benchmarks for 120 saddle-searches from Fe SIA configuration.

Search method	Dimer	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local	Local
Max steps in init. direction	0	0	3	10
Success rate (%)	87.5	98.3	98.3	98.3
Execution time (s)	401	54	48	51
Saddles at 0.30-0.32 eV	40	99	91	101
Saddles at 0.41-0.42 eV	5	16	27	16
Saddles at 0.68 eV	4	2	0	0
Other saddles	56	1	0	1

Table 5: Benchmarks for 600 saddle-searches from Fe 50-vacancy configuration.

Search method	Dimer	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local	Local
Max steps in init. direction	0	0	5	20
Success rate (%)	100	100	100	99.5
Execution time (s)	166	126	156	260

Table 6: Benchmarks for 144 saddle-searches from an Fe configuration with two interlocked dislocation loops.

Search method	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local
Max steps in init. direction	0	3	20
Success rate (%)	32.6	46.5	29.9
Execution time (s)	339	258	349

Table 7: Benchmarks for 144 saddle-searches from an FeCr configuration with two interlocked dislocation loops.

Search method	Dimer	ARTn	ARTn
Initial displacement	Global	Local	Local
Max steps in init. direction	0	3	20
Success rate (%)	5.6	32.6	29.9
Execution time (s)	2248	383	381